

NVM Express Technical Errata

Errata ID	026
Change Date	3/7/2012
Affected Spec Ver.	NVM Express 1.0c
Corrected Spec Ver.	

Submission info

Name	Company	Date
Judy Brock	Samsung	3/7/2012
Santosh Singh	Samsung	3/7/2012
Amber Huffman	Intel	3/7/2012

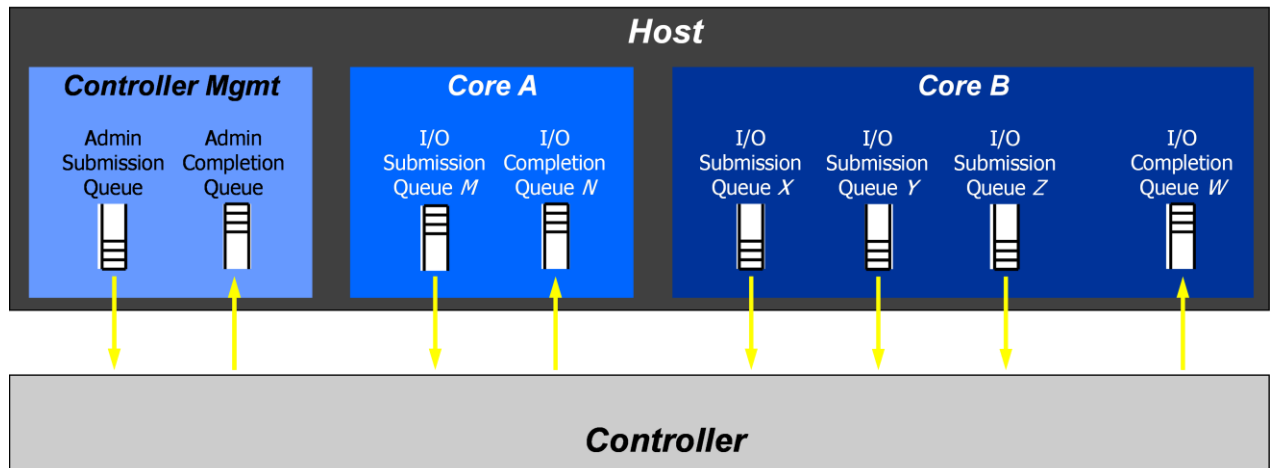
This erratum clarifies the introduction material on the queue pair model for the Admin Queue. Specifically, the Admin Queue is always a one-to-one pairing between the Admin Submission Queue and Admin Completion Queue regardless of the queue pair model used by IO Submission and Completion Queues created by host software.

This erratum fixes two typos in the specification.

Description of the specification technical flaw:

Replace Figure 2 with the updated figure below that includes the Controller Management box:

Figure 2: Queue Pair Example, $n:1$ Mapping



Update the paragraph preceding Figure 1 as shown:

Host software creates queues, up to the maximum supported by the controller. Typically the number of command queues created is based on the system configuration and anticipated workload. For example, on a four core processor based system, there may be a queue pair per core to avoid locking and ensure data structures are created in the appropriate processor core's cache. Figure 1 provides a graphical representation of the queue pair mechanism, showing a 1:1 mapping between Submission Queues and Completion Queues. Figure 2 shows an example where multiple I/O Submission Queues utilize the same I/O Completion Queue on Core B. **Figure 1 and Figure 2 show that there is always a 1:1 mapping between the Admin Submission Queue and Admin Completion Queue.**

Update the second to last paragraph of section 8.4 as shown:

The default NVM Express power state is implementation specific and shall correspond to a state that does not consume more power than the lowest value specified in the form factor specification used by the PCI Express SSD. The host shall never select a power state that consumes more power than the PCI Express slot power limit control value expressed by the Captured Slot Power Limit Value (~~CSPLC~~) (**CSPLV**) and Captured Slot Power Limit Scale (CSPLS) fields of the PCI Express Device Capabilities (PXDCAP) register. Hosts that do not dynamically manage power should set the power state to the lowest numbered state that satisfies the PCI Express slot power limit control value.

Modify Figure 73 as shown below:

Figure 73: Set Features – Feature Identifiers

Feature Identifier	O/M	Persistent Across Power States	Uses Memory Buffer for Attributes	Description
00h				Reserved
01h	M	No	No	Arbitration
02h	M	No	No	Power Management
03h	M O	Yes	Yes	LBA Range Type
04h	M	No	No	Temperature Threshold
05h	M	No	No	Error Recovery
06h	O	No	No	Volatile Write Cache
07h	M	No	No	Number of Queues
08h	M	No	No	Interrupt Coalescing
09h	M	No	No	Interrupt Vector Configuration
0Ah	M	No	No	Write Atomicity
0Bh	M	No	No	Asynchronous Event Configuration
0Ch – 7Fh				Reserved
80h – BFh				Command Set Specific (Reserved)
C0h – FFh				Vendor Specific

O/M: O = Optional, M = Mandatory

Disposition log

3/7/2012	Erratum captured.
4/18/2012	Erratum ratified.

Technical input submitted to the NVMHCI Workgroup is subject to the terms of the NVMHCI Contributor's agreement.